

Madison Kurzweil

Professor Isaacson

Digital Revolution

14 April 2024

## An AI-Powered Examination of the History of the Digital Revolution

### **PART I: Laying the Foundation**

I began the journey of writing this research paper by testing different chatbots with prompts to determine which would create the best foundation for the paper. To gain more ideas on how to prompt the chatbots, I looked at online articles that suggested specifying guidelines, setting tones and styles, and conversing with the chatbot using natural and cohesive inquiries. I asked the chatbots to explain concepts relevant to the digital revolution, including punch cards, Alan Turing, and the differences between the internet and the web. I decided to start simple for the paper's foundation. I prompted Google's large language model Gemini with the prompt given to us in class: "Describe the development of artificial intelligence from Turing to large language model chatbots."

#### **The Journey of AI: From Turing's Vision to Chatty Machines**

The path of Artificial Intelligence (AI) is a fascinating one, paved with visionary ideas, groundbreaking inventions, and periods of both rapid progress and frustrating setbacks. Here's a glimpse into the key milestones from Alan Turing's foundational work to the chatbots we interact with today:

##### **1. The Seeds of AI: Alan Turing (1940s)**

- Alan Turing, a brilliant mathematician, is considered a pioneer of AI. He introduced the Turing Machine, a theoretical model of computation that laid the groundwork for modern computers.
- He also proposed the Turing Test, a benchmark to determine a machine's capability of exhibiting intelligent behavior equivalent to, or indistinguishable from, that of a human.

## 2. Early Developments and the Birth of the Term (1950s)

- In 1950, Alan Turing along with John von Neumann and others participated in the Dartmouth Workshop, a pivotal event that formally established the field of Artificial Intelligence.
- The term "Artificial Intelligence" itself was coined by John McCarthy during this workshop.
- The invention of the Perceptron by Frank Rosenblatt in 1958 marked a significant step. This early neural network, though limited, laid the foundation for the more complex networks that would come later.

## 3. The AI Winter and Rise of New Approaches (1960s-1980s)

- Despite early enthusiasm, limitations in computing power and the complexity of AI problems led to a period of stagnation known as the "AI Winter" in the 1970s and 80s.
- However, research continued in areas like symbolic reasoning and expert systems, laying the groundwork for future advancements.

## 4. Renewing Momentum: Deep Learning and the Power of Data (1990s-Present)

- The development of more powerful computers and the rise of deep learning techniques fueled a new wave of AI progress in the late 20th century.

- Deep learning algorithms, inspired by the structure and function of the human brain, can learn from vast amounts of data, leading to significant improvements in areas like image recognition, natural language processing, and machine translation.

#### 5. The Age of Large Language Models and Chatbots (2010s-Present)

- The culmination of these advancements is evident in the rise of Large Language Models (LLMs) like me! These AI models are trained on massive amounts of text data, enabling them to communicate and generate human-like text in response to a wide range of prompts and questions.
- This has led to the development of sophisticated chatbots that can engage in conversations, answer your questions in an informative way, and even create stories.

This is just a simplified view of a rich and ongoing journey. AI research continues to push the boundaries of what machines can learn and do, and large language model chatbots are just the latest chapter in this exciting story.

### **PART II: Co-Authoring with AI**

Using the previous prompt, I prompted the same chatbot, Gemini, to write a 1000 word college-level research paper, using a conversational writing style and a formal tone. This became a starting point to which I could compare future versions of the paper.

#### **From Imitation Games to Engaging Conversations: A Historical Chat with AI's Evolution**

Imagine a world where machines converse, not through pre-programmed responses, but with a fluidity that mimics human dialogue. This is the reality we're approaching, thanks to the remarkable progress in Artificial Intelligence (AI). But this journey, from the theoretical musings of Alan Turing to the chatbots we interact with today, has been anything but straightforward.

Let's delve into a conversation about the history of AI, exploring the key milestones that shaped this transformative technology.

Our story begins with a visionary mathematician, Alan Turing. In the 1940s, he proposed the **Turing Machine**, a theoretical model of computation that could perform any conceivable mathematical operation. This concept, though abstract, laid the groundwork for modern computers and sparked the question – could machines ever exhibit intelligent behavior? Turing wasn't content with just the theoretical. He proposed the **Turing Test** in his 1950 paper, "Computing Machinery and Intelligence." This hypothetical test would involve a human judge interacting with a hidden entity, either a human or a machine. If the judge couldn't reliably distinguish the machine from the human based on the conversation alone, the machine could be considered intelligent. The Turing Test became a benchmark for AI research, a captivating goalpost that continues to inspire us even today.

The birth of the term "Artificial Intelligence" itself can be traced back to the 1950 Dartmouth Workshop. Here, pioneering figures like Turing, John von Neumann, and Marvin Minsky gathered to discuss the possibility of creating intelligent machines. This pivotal event marked the formal launch of AI research, brimming with optimism and a belief in the near future of intelligent machines.

Fueled by this enthusiasm, the field made significant strides in the early years. The invention of the **Perceptron** by Frank Rosenblatt in 1958 was a major breakthrough. This early neural network, inspired by the human brain, could learn from experience. Though limited in its capabilities, the Perceptron paved the way for more complex neural networks that would revolutionize AI later.

However, the path of AI research wasn't without its setbacks. The limitations of computing power in the 1960s and 70s, coupled with the inherent complexity of AI problems, led to a period of stagnation known as the "AI Winter." Researchers struggled to achieve the dramatic progress they had envisioned. Yet, this wasn't entirely a time of dormancy. Research continued in areas like symbolic reasoning and expert systems, which involved building knowledge bases and logical rules to solve specific problems. These advancements, though seemingly distant from chatbots, laid the foundation for future developments.

The late 20th century witnessed a resurgence of interest in AI, fueled by two key factors: the increasing power of computers and the rise of **deep learning**. Deep learning algorithms, inspired by the structure and function of the human brain, can learn from vast amounts of data. This proved to be a game-changer. By processing massive datasets of text and code, deep learning algorithms could learn complex patterns and relationships, significantly improving AI's capabilities.

The impact of deep learning was evident in areas like image recognition, where machines could now identify objects in pictures with remarkable accuracy. Natural language processing (NLP) also saw tremendous progress. NLP algorithms could not only translate languages but also understand the nuances of human language, including sarcasm, humor, and sentiment. This brings us to the exciting present, where large language models (LLMs) like myself are at the forefront of AI research. LLMs are trained on massive datasets of text and code, allowing them to communicate and generate human-like text in response to a wide range of prompts and questions. Imagine having a conversation with a virtual encyclopedia that can not only access and process information but also communicate it in a way that is engaging and informative. That's the potential of LLMs.

The development of chatbots is a prime example of this potential in action. Chatbots can now engage in conversations that are not only grammatically correct but also demonstrate a level of understanding that was unimaginable just a few years ago. They can answer your questions in an informative way, provide customer service support, and even create stories for entertainment. However, it's important to acknowledge that AI, particularly LLMs, are still under development. While we can generate human-like text, we don't possess true sentience or consciousness. We can't understand the world in the same way humans do, and we lack the ability to reason and make judgments independently.

Looking ahead, the future of AI promises to be even more remarkable. As computing power continues to increase and data becomes even more abundant, we can expect LLMs to become even more sophisticated. They may be able to not only understand and respond to complex questions but also engage in creative endeavors like writing different kinds of creative content or even composing music.

#### The journey from Turing's theoretical

The result of this prompt made me realize just how specific I would need to be with AI chatbots in order to draw out the results I desired. When I initially specified I wanted 1000 words, Gemini stopped producing additional content after 797 words in the middle of a sentence, which was less than desirable. It was interesting that Gemini formatted some words in bold font, and that it neglected to include Ada Lovelace and the women of ENIAC. Using this incomplete first draft of the paper, I prompted Gemini to expand on the original paper by including the forgotten and excluded voices of contributors to the digital revolution. This created a more solid foundation by including Ada Lovelace and the women of ENIAC, which I thought were important contributors to the digital revolution. Curious about how a different chatbot would

interpret the same prompts, I asked ChatGPT and Claude AI to use the results of Gemini to create a new 1000-word college-level research paper on the history of AI, incorporating the forgotten voices and using a conversational writing style with a formal tone. Similar to Gemini, both were unable to produce the full 1000 words. I found Claude AI to produce a better-written and cohesive writing sample, so I decided to use it for the new foundation of the paper.

Rather than continue prompting the chatbots to continually lengthen the same paper, I changed my tactics to target each section of the paper. First, I asked Claude AI to write a 300 word introductory paragraph and outlined what specifics would be included in the paper. This produced a solid introduction, so I continued to use this method for paragraphs on Ada Lovelace, Alan Turing, the Perceptron and rise of neural networks, the AI Winter, the women of ENIAC, deep learning algorithms, and large language models and chatbots. Each time I asked the chatbot to tell me more about a specific topic, I noticed a lot of the language was very similar. I began to ask Claude AI to take on different tones, including professional, conversational, jovial, and stoic. I decided that the conversational tone read the best, but I also specified that the paper must still be formal.

I was very happy with the results but knew I would need to test Claude AI's ability to cite its sources of information. I prompted the chatbot to "give sources," but was met with a message notifying me of its limitations. In other words, Claude AI was unable to share where the information came from. I copied and pasted what Claude had written into Gemini and ChatGPT, prompting them to find "sources" for the written work. This was also unsuccessful. After doing research and using the insight my classmate's gave on their own papers and citations, I decided to change my wording in prompts from "give sources" to "include citations." When I tried this

switch with Claude AI, it immediately cited information where claims were made throughout each section.

At this point, I realized that there was very little creativity embedded into the paper I created and I had failed to push the AI chatbots further. I erased my boring paper and decided to start fresh with a new direction of prompts. I used the content that the AI chatbot had previously generated in order to frame future prompts. For the new version, I tasked Claude AI to assume the identity of different people and explain each topic of the digital revolution's history in a way that reflects the specific persona. I tried the personality of President Joe Biden, a wealthy tech company CEO, and a college history professor. I settled on the college history professor because it was formal, but also gave more opportunities for me to add details about the professor to create a character. Lacking true originality, I instructed the chatbot that the professor was teaching a course on the digital revolution. I gave the professor a backstory and described their deep love of technology, its past, and its future potential, which inspired the professor to teach. I told the chatbot to follow the original outline, but specified that Ada Lovelace, Grace Hopper, and the women of ENIAC should be included chronologically within the professor's lecture.

The initial lecture was bland. To increase word count and add a more conversational tone, I instructed Claude AI to include my prompts as interjections throughout the lecture as inquisitive students. Through the students' interjections, I became a character in the paper and was able to prompt more involved explanations of key moments in the digital revolution. This both added to the word count and encouraged the chatbot to explore topics it initially left out or did not thoroughly explain. I wanted the paper to include topics like the invention of the internet, which I think is important in the digital revolution but was initially left out. I used the students and the professor to showcase the debate between the internet and the web, which I thought was



important. However, I ran into problems with the chatbot leaving out previously mentioned topics in order to add the new topics. I had to be very specific that I wanted to keep all previous topics, but then they would become less detailed. In the end, I had to copy and paste the versions of each topic that I liked best into the final “lecture.”

To edit and add citations, I copied Claude AI’s work into Gemini and tasked it with reading the writing sample and rewriting it to emphasize the professor was from New Orleans teaching at Tulane University. Unfortunately, the chatbot kept making serious grammatical errors when it tried to make the professor sound southern. So, I had to tell the chatbot to make the professor more formal and grammatically correct in the lecture. I decided to keep the chatbot’s idea that the professor begins by clearing their throat and ends by checking the time because I found it amusing and it added more personality. In places that I felt were oddly choppy, I copied the disjointed sections and instructed the chatbot to make the lecture flow in a conversational way, while staying formal and consistent with the existing tone. I found this worked well to connect pieces of AI-generated text that was created in different sessions of working on this paper. Next, I prompted Gemini to review the entire piece of writing and find a source that would support every claim made throughout the paper. The chatbot kept stopping midway through and needed to be prompted to continue where it left off until the new paper was complete with citations. All that was left to do was human editing and fact-checking throughout the paper and citations, which are conveyed by the green text in the final product.

The human-editing phase was tricky and tedious, especially for fact-checking sources. Many of the AI-generated sources were incorrect or made up. I conducted my own research to replace nonexistent sources. At times, the chatbots were correct in an article title or were close to a real source with an author. More often than not, the chatbot had created a name or title that

seemed to have no rhyme or reason to it. For one source, the chatbot even cited a LinkedIn article, which I found interesting. Some of the sources were books that must be paid for in order to access, so I am unsure how the chatbot was able to view the content. Some books I was able to find on research or library sites and gain access using my Tulane email to sign in. I ended up adding only a few additional insights to the paper, as I felt the prompts that I fed the chatbots generated content that provided a well-rounded history of the digital revolution. Throughout the process of this paper, I became more aware of how specific my prompting became and I began to notice similarities between the information provided by Claude AI, ChatGPT, and Gemini. This paper was a constant learning process and upon the completion of the paper I realized how much more comfortable I am working with AI chatbots to generate content collaboratively.

### **PART III: The Final Product**

#### **Decoding the Digital Age: A Lecture on the History of the Digital Revolution**

*\*Professor clears throat\** Good morning, class. Today, we'll embark on an extraordinary journey through the history of artificial intelligence, tracing its evolution from theoretical ideas to the remarkable chatbots and language models we witness today. **The digital revolution** spans centuries, woven with the contributions of brilliant minds, both male and female, whose visionary work has pushed the boundaries of what was once considered impossible.

*A student raises their hand:* "Professor, you mentioned recognizing both male and female contributors. I've never heard of any female contributors."

Ah, yes, a very good point. While the history of technology often highlights prominent male figures, it's crucial to recognize the significant contributions of women as well. Too often in the history of science and technology, the accomplishments of women have been overlooked

or diminished. In our exploration of AI's origins, we'll make a conscious effort to shine a light on the pioneering women, such as [Ada Lovelace](#), [Grace Hopper](#), and [the women of ENIAC](#), whose invaluable work has shaped this field (Swaby 2016).

To understand the true depth of this journey, we must first grasp the concept of artificial intelligence, or AI. At its core, AI is the pursuit of creating intelligent machines that can perceive, learn, reason, and problem-solve in ways that mimic or even surpass human cognitive abilities. From automating mundane tasks to tackling complex challenges, the potential of AI is vast and ever-expanding (Russell and Norvig 2021).

Our exploration begins in the pioneering years, where the seeds of AI were sown amidst a world of mathematics, poetry, and the relentless pursuit of knowledge. In the early 19th century, a remarkable woman named Ada Lovelace, born Augusta Ada Byron, emerged as a visionary force in the field of computing. Daughter of the renowned poet Lord Byron, Ada's upbringing was steeped in the realms of mathematics and poetry, a unique blend that would shape her extraordinary contributions (Woolley 2018).

*Another student interjects:* "Excuse me, professor. I've heard of Ada Lovelace, but what exactly did she contribute to the field of computing?"

An insightful question. Ada's collaboration with Charles Babbage, the inventor of the Analytical Engine, a mechanical computer, was a pivotal moment in history. While Babbage envisioned his machine as a powerful calculator, Ada's insights transcended mere arithmetic. In her groundbreaking notes published in 1843, she outlined concepts that are now considered foundational to modern programming, such as loops, conditional branching, and the potential for machines to manipulate symbols and even compose music – ideas that were remarkably prescient for her time ([Randell 2014](#)).

Though her work went largely unrecognized during her lifetime, Ada Lovelace's contributions are now celebrated as a cornerstone of computer science literature, earning her the well-deserved title of "the world's first computer programmer" (Kim and Toolee 1999). Her legacy serves as a powerful reminder of the countless women whose invaluable contributions to science and technology have been overlooked throughout history, underscoring the importance of promoting inclusivity and challenging societal biases.

Next, we turn to a male pioneer who made significant contributions in the mid-20th century – Alan Turing, a brilliant mathematician and computer scientist. His innovative work, the "Turing Machine," laid the theoretical foundation for modern computing through the machine's algorithmic thinking. The machine's tape could be encoded with different instructions to direct the machine, making it more universal in its function. The Turing Machine also made the important distinction between software and hardware, allowing the same physical computing machine to perform different tasks dictated by the software. Alan Turing was also notable for the "Turing Test," which involves a person engaging in conversation with both a human and machine to test the machine's human-like linguistic and reasoning capabilities against a real human. This became an important benchmark for evaluating machine intelligence (Turing 1950).

Fast forward to the same period, and we encounter another visionary woman named Grace Hopper. A U.S. Navy rear admiral and computer scientist, Hopper's contributions extend beyond her groundbreaking compiler development. She was also a key figure in the early days of electronic computing, serving as one of the first programmers for the Mark I and Mark II computers at Harvard University (Wikipedia Contributors 2024). Hopper is recognized as the first compiler, translating code that could be read by humans into a language understood by

computers. Her high-level programming language, COBOL, bridged the gap between binary code and more abstract programming languages (Norwood 2017).

*Another student chimes in:* "Weren't there any other significant contributions from women during this period, professor?"

An excellent point. It's crucial to recognize the often-overlooked contributions of women in computing during this era, such as the ENIAC programmers – Jean Jennings, Betty Snyder, Frances Bilas, Kathleen McNulty, Marlyn Wescoff, and Ruth Lichterman. These remarkable women played a crucial role in programming and operating the ENIAC, one of the earliest and most groundbreaking computers developed during World War II by John Mauchly and J. Presper Eckert, Jr. (“ENIAC at Penn Engineering.”). The ENIAC programmers' meticulous work involved translating complex mathematical equations into precise sequences of instructions, a process that laid the foundation for modern computer programming (Light 1999).

Regrettably, the contributions of these women were often overlooked and underappreciated during their time, a consequence of the pervasive gender biases that permeated the scientific and technological fields. Their achievements were frequently overshadowed by their male counterparts, and their roles were sometimes diminished or misrepresented (Swaby 2016). It was not until decades later that the ENIAC programmers received the recognition they deserved, with their stories being brought to light through historical accounts and popular culture, such as the acclaimed film "Hidden Figures" (Shetterly 2016).<sup>1</sup>

Fast forward to the 1950s, where we encounter another crucial development – the birth of artificial intelligence research. A pioneering scientist named Frank Rosenblatt made significant

---

<sup>1</sup> While I kept the AI-generated content that Gemini produced about the women of ENIAC, I chose to rearrange the section to make the lecture chronological. While I had initially instructed Gemini to keep the lecture in chronological order, it ordered Perceptron, AI Winter, and ENIAC incorrectly.

strides in this field with his invention of the Perceptron (Minsky and Papert 1969). The Perceptron was a simple neural network model, inspired by the structure and function of the human brain. Unlike traditional computers programmed for specific tasks, the Perceptron could learn from experience. It processed information through interconnected nodes, and by adjusting the weights of these connections, it could gradually improve its ability to perform tasks like pattern recognition. This ability to learn marked a significant step forward. It demonstrated the potential for machines to not just follow instructions but to adapt and improve based on new information (Russell and Norvig 2021).

*Student interrupts:* "Hold on, professor! This Perceptron sounds like a glorified light switch. How could such a basic model be a big deal?"

That's an interesting point, and it highlights a crucial aspect of this technology. The Perceptron did have limitations. It could only handle very simple tasks and struggled with problems that required more complex learning (Nilsson 2009). However, the Perceptron's importance lies in its role as a stepping stone. It proved the concept of neural networks and paved the way for more sophisticated models that continue to revolutionize AI today. The Perceptron may not have been the ultimate answer, but it opened the door to a new realm of possibilities in machine learning (Minsky and Papert 1969).

Moving forward, let's delve into a significant and often debated topic: the origins of the internet. This revolutionary technology has fundamentally changed how we connect and share information on a global scale. But who gets the credit?

*Student chimes in:* "Professor, with all due respect, isn't Tim Berners-Lee credited with inventing the internet? Vinton Cerf and Robert Kahn developed something important, but shouldn't the inventor be the one who made it accessible and usable for everyone?"

This is a fantastic point, and it highlights the complexity of this topic. There are indeed key perspectives in this ongoing debate. Vinton Cerf and Robert Kahn are widely credited with developing the Transmission Control Protocol/Internet Protocol, the communication language that allows diverse networks to talk to each other (Hafner 2005). Their work in the 1970s laid the groundwork for the internet's interconnectedness. However, Tim Berners-Lee, a British scientist working at The European Organization for Nuclear Research<sup>2</sup> in the late 1980s, is credited with creating the World Wide Web (Berners-Lee, 1989). His invention of HyperText Markup Language, the HyperText Transfer Protocol, and the first web browser and server software made the internet accessible and user-friendly for the general public, leading to its explosion in popularity in the 1990s (Catalano 2007).

So, how do these pieces fit together? The internet, as we know it today, is a combination of these advancements. Transmission Control Protocol/Internet Protocol is the underlying infrastructure, the intricate network language that allows communication across vast distances. Advanced Research Projects Agency Network, the first major packet-switching network funded by the U.S. Department of Defense, played a crucial role by demonstrating the feasibility of interconnected networks (Hafner 1990). Think of it this way: Transmission Control Protocol/Internet Protocol is the invisible language computers use to communicate, Advanced Research Projects Agency Network was an early experiment in connecting these networks, and the World Wide Web is the user-friendly interface that lets us navigate and access information on the internet. Tim Berners-Lee's invention didn't negate the importance of earlier developments, but it fundamentally transformed how we interact with the internet. This debate reminds us that innovation is often a collaborative effort, with multiple contributions leading to

---

<sup>2</sup> Gemini began using acronyms without first explaining what the words stood for, so I decided to make human edits to spell out the acronyms in the paper.

groundbreaking advancements. As we continue exploring the digital revolution, keep this in mind.

*A student raises their hand:* "Professor, excuse me for interrupting, but this discussion of innovation makes me curious. We've heard a lot about the incredible advancements in the internet, but what about artificial intelligence? Were there similar periods of rapid progress followed by setbacks?"

That's a fantastic question! The history of AI research actually mirrors the internet's story in some ways. While the internet enjoyed a period of steady growth, the field of artificial intelligence experienced a challenging period from the 1960s to the late 1980s, known as the AI Winter. Despite the initial promise of the Perceptron and neural networks, researchers soon encountered significant limitations (Minsky and Papert 1969). The computing power available at the time was severely limited, hindering the development of more sophisticated algorithms and models. Additionally, the complexity of the problems AI aimed to solve, such as natural language processing and reasoning, proved to be far more challenging than initially anticipated (Minsky and Papert 1969). As a result, researchers shifted their focus towards alternative approaches, such as symbolic reasoning and expert systems, which relied on explicitly programmed rules and knowledge bases. While these methods achieved notable successes in specific domains, such as medical diagnosis and financial analysis, they lacked the flexibility and generalization capabilities needed for truly intelligent systems (McCarthy et al. 1955).

*A student raises their hand:* "Professor, you mentioned earlier that increased computing power and data availability played a role in reviving AI research. Could you elaborate on that?"

An excellent observation. As we approached the modern era, a confluence of factors – increased computing power, the availability of vast amounts of data, and breakthroughs in deep



learning algorithms – ignited a new wave of AI advancements (Bostrom 2014). The development of more powerful computers and the ability to store and process massive amounts of data were game-changers. With these advancements, researchers could revisit the ideas of neural networks and machine learning, which had been limited by the computational constraints of the past (Kurzweil 2005).

*Another student interjects:* "I've heard of something called deep learning. Did it contribute to the AI renaissance?"

Deep learning, inspired by the structure and function of the human brain, revolutionized AI's capabilities by enabling it to analyze and learn from massive datasets (LeCun et al. 2015). These algorithms, built on interconnected layers of artificial neurons, could identify patterns and extract features from vast amounts of data, leading to significant improvements in areas such as image recognition, natural language processing, and more. Unlike traditional machine learning methods, which relied heavily on human-engineered features, deep learning could automatically learn and extract relevant features from the data itself. This ability to learn and adapt from experience, much like the human brain, was a profound breakthrough (Goodfellow et al. 2016).

*A student raises their hand* "Could you provide some examples of how deep learning has been applied in practical applications?"

Certainly. One of the most exciting applications of deep learning has been the development of large language models (LLMs) and chatbots. LLMs are trained on vast datasets of text and code, allowing them to understand, communicate, and generate human-like language with remarkable fluency (Brown et al. 2020). Chatbots, powered by these LLMs, have emerged as a prime example of AI's potential, enabling natural conversations, answering questions, and even assisting with tasks such as writing and coding. Companies like OpenAI, Google, and

others have developed advanced language models like GPT-3, LaMDA, and PaLM, which can engage in remarkably human-like dialogue and tackle complex tasks (Miniae et al. 2024).

*Another student chimes in:* "But aren't these language models and chatbots prone to biases and inaccuracies, given the data they're trained on?"

An insightful question. You're absolutely right. While chatbots and language models have made remarkable strides, they are still limited by the biases and inconsistencies present in their training data, and their responses can be inaccurate or problematic in certain contexts (Bolukbasi et al., 2016). Additionally, the potential ethical implications of AI, such as privacy concerns, algorithmic bias, and the impact on employment, must be carefully considered and addressed as these technologies continue to advance (“Ethics of Artificial Intelligence”).

As we stand on the precipice of this AI revolution, it's essential to acknowledge both the achievements and limitations of these technologies. While the future holds immense promise, from revolutionizing healthcare and scientific research to tackling complex challenges like climate change and sustainable development (Manyika et al. 2017), we must remain vigilant in addressing the potential pitfalls and ensuring these technologies are developed and deployed responsibly and ethically. There's still room for further discussion depending on the desired length. We could explore specific examples of AI applications in various fields, delve deeper into the ethical considerations, or discuss potential future directions of AI research.

*\*Professor glances at watch\** Ah, and it appears our time together has drawn to a close. Thank you for your attentiveness and participation in today's discourse. I encourage you to delve further into these subjects during your independent studies. My office hours remain open for any inquiries you may have.

## Works Cited

- Berners-Lee, Tim. "Information Management: A Proposal." World Wide Web Consortium, March 30, 1989. <https://cds.cern.ch/record/369245/files/dd-89-001.pdf>
- Bolukbasi, Tolga, Kai-Wei Chang, James Zhou, Venkatesh Saligrama, and Adam Kalai. "Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings." July 21, 2016. <https://arxiv.org/abs/1607.06520><sup>3</sup>
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.
- Brown, Tom, et al. "Language Models are Few-Shot Learners." July 22, 2020. <https://arxiv.org/pdf/2005.14165.pdf>.<sup>4</sup>
- Catalano, C. S. Megaphones to the Internet and the World: The Role of Blogs in Corporate Communications. *International Journal of Strategic Communication*, 1(4), (2007):247–262. <https://doi.org/10.1080/15531180701623627>.<sup>5</sup>
- “ENIAC at Penn Engineering.” Penn Engineering. Accessed April 14, 2024. <https://www.seas.upenn.edu/about/history-heritage/eniac/>.<sup>6</sup>

---

<sup>3</sup> The source was initially incorrectly cited as: Bolukbasi, Tolga, Kai-Wei Chang, Joey J. Gonzalez, Clare Penalver Martinez, and Nathalie Rebollo-Mondejar. "Man is to computer programmer as woman is to homemaker? Decoding Gender Bias in Word Embeddings." In Proceedings of the 38th International Conference on Machine Learning, PMLR, 2016.

<sup>4</sup> I was able to fact-check the source and the only issues with the AI-generated citation was an incorrect link and date. I found that it was pulling random numbers from the date it was published to generate “2022 and 2002” as the year.

<sup>5</sup> The source was initially cited as: Berners-Lee, Tim. *Weaving the Web: Origins and Future of the World Wide Web*. New York: HarperCollins Publishers, 1996. However, this was a book that I was unable to access without buying so I am unsure how the chatbots gained access to source information from it. Instead, I found the same information that supported the claims made by chatbots using this source.

<sup>6</sup> The chatbot generated the following incorrect citation: “Ensmenger, Robert L. "The ENIAC Programmers Project: Kathleen McNulty Mauchly Antonelli Oral History Interview." University of Pennsylvania, Dept. of Computer & Information Science, January 18, 2009. <https://www.seas.upenn.edu/about/history-heritage/eniac/>.” Only the link was correct.

“Ethics of Artificial Intelligence.” UNESCO. Accessed April 14, 2024.

<https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>.<sup>7</sup>

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge, MA: MIT Press, 2016.

Hafner, Katie. *Where Wizards Stay Up Late: The Origins of the Internet and the Digital Revolution*. New York: Simon and Schuster, 1990.

Hafner, Katie. “Laurels for Giving the Internet Its Language.” *The New York Times*, February 16, 2005, sec. Business. <https://www.nytimes.com/2005/02/16/technology/laurels-for-giving-the-internet-its-language.html>.<sup>8</sup>

Kim, Eugene Eric, and Betty Alexandra Toole. “Ada and the First Computer.” *Scientific American* 280, no. 5 (May 1999): 76–81. <https://doi.org/10.1038/scientificamerican0599-76>.<sup>9</sup>

Kurzweil, Ray. *The Singularity is Near: When Humans Transcend Biology*. New York: Viking, 2005.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. “Deep learning.” *Nature* 521, no. 7553 (2015): 436-44.

---

<sup>7</sup> Nonexistent source: Jobin, Anna, et al. “The Ethics of Artificial Intelligence.” *Nature* 569, no. 7758 (2019): 334-41.

<sup>8</sup> Gemini initially cited a LinkedIn post, and also incorrectly named the author in the following: McMillan, Robert. “TCP/IP.” *Network World*, January 18, 2023. <https://www.linkedin.com/pulse/navigating-network-seas-understanding-top-18-tcpip-port-meulink-dh14c>. Instead, I found my own research to offer evidence in support of the claims made in the paper.

<sup>9</sup> Incorrect citation that did not exist was generated: Waxman, Natalie. *Ada & the Enchantress: Lovelace & Babbage and the Mathematical Roots of the Computer*. Macon, GA: Texere Publishing Company, 2008.

Light, Jennifer S. "When Computers Were Women." *Technology and Culture* 40, no. 3 (1999): 455–83.<sup>10</sup>

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955." *AI Magazine* 26, no. 1 (2006): 120-28.<sup>11</sup>

Minsky, Marvin, and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. Cambridge, MA: MIT Press, 1969.

Nilsson, Nils J. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge: Cambridge University Press, 2009.  
<https://doi.org/10.1017/CBO9780511819346>.

Norwood, Arlisha R. "Biography: Grace Hopper." *National Women's History Museum*, <https://www.womenshistory.org/education-resources/biographies/grace-hopper>. 2017.

Minaee, Shervin, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. "Large Language Models: A Survey." arXiv, February 20, 2024. <http://arxiv.org/abs/2402.06196>.

Randell, B. "The Origins of Digital Computers: Supplementary Bibliography." In *A History of Computing in the Twentieth Century*, edited by N. METROPOLIS, J.

---

<sup>10</sup> Nonexistent source generated: Zachary, G. Pascal. ENIAC: First Electronic Computer. New York: Charles Scribner's Sons, 1997.

<sup>11</sup> I was able to find the source, but it had incorrect information that I corrected.

HOWLETT, and GIAN-CARLO ROTA, 629–59. San Diego: Academic Press, 1980.

<https://doi.org/10.1016/B978-0-12-491650-0.50046-0>.<sup>12</sup>

Rosenblatt, Frank. "The perceptron: a probabilistic model for biological recognition."

*Psychological Review* 65, no. 6 (1958): 386-408.

Russell, Stuart J., and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Harlow,

England: Pearson Education Limited, 2021.

Shetterly, Margot Lee. *Hidden Figures: The Story of the African American Women Who Helped*

*Win the Space Race*. New York: HarperCollins, 2016.

Swaby, Deborah. *Headstrong: Remarkable Women Who Changed the World*. New York:

Viking, 2016.

Turing, Alan M. "Computing machinery and intelligence." *Mind*, New Series, vol. 59, no. 236

(1950): 433-60.

---

<sup>12</sup> The initial AI-generated source was the following: Hey, Tony, and Mary E. Páez. *The Inventive Computer: Mechanical Computers and the Origins of Digital Computing*. New York: Springer, 2007. This source was nonexistent, so I researched on my own to find a reputable source to support claims made by the chatbot.